Noname manuscript No.
(will be inserted by the editor)

Sketch-based Image Retrieval from Millions of Images under Rotation, Translation and Scale Variations

Sarthak Parui · Anurag Mittal

arXiv:1511.00099v1 [cs.CV] 31 Oct 2015

Received: date / Accepted: date

Abstract Proliferation of touch-based devices has made sketch-based image retrieval practical. While many methods exist for sketch-based object detection/image retrieval on small datasets, relatively less work has been done on large (web)-scale image retrieval. In this paper, we present an efficient approach for image retrieval from millions of images based on user-drawn sketches. Unlike existing methods for this problem which are sensitive to even translation or scale variations, our method handles rotation, translation, scale (i.e. a similarity transformation) and small deformations. The object boundaries are represented as chains of connected segments and the database images are pre-processed to obtain such chains that have a high chance of containing the object. This is accomplished using two approaches in this work: a) extracting long chains in contour segment networks and b) extracting boundaries of segmented object proposals. These chains are then represented by similarity-invariant variable length descriptors. Descriptor similarities are computed by a fast Dynamic Programming-based partial matching algorithm. This matching mechanism is used to generate a hierarchical k-medoids based indexing structure for the extracted chains of all database images in an offline process which is used to efficiently retrieve a small set of possible matched images for query chains. Finally, a geometric verification step is employed to test geometric consistency of multiple chain matches to improve results. Qualitative and quantitative results clearly demonstrate superiority of the approach over existing methods.

Keywords Sketch-based Retrieval · Image Retrieval · Shape Representation · Indexing

Sarthak Parui · Anurag Mittal
Computer Vision Lab,
Dept. of Computer Science & Engineering
Indian Institute of Technology Madras
Tel.: +91-44-22575352
E-mail: sarthak,amittal@cse.iitm.ac.in

1 Introduction

The explosive growth of digital images on the web has substantially increased the need for an accurate, efficient and user-friendly large-scale image retrieval system. With the growing popularity of touch-based smart computing devices and the consequent ease and simplicity of querying images via hand-drawn sketches on touch screens (Lee *et al.*, 2011), sketch-based image retrieval has emerged as an interesting application. The standard mechanism of text-based querying could be imprecise due to wide demographic variations. It also faces the issue of availability, authenticity and ambiguity in the tag and text information surrounding an image (Sigurbjörnsson and Van Zwol, 2008; Schroff *et al.*, 2011), which necessitates us to exploit image content for better search. Although various popular image search engines such as Google[1] and TinEye[2] provide an interface for similar image search using an exemplar image as the query, a user may not have access to such an image every time. Instead, a hand-drawn sketch may be used for querying since sketching is a fundamental mechanism for humans to conceptualize and render visual content as also suggested by various Neuroscience studies (Marr, 1982; Landay and Myers, 2001; Walther *et al.*, 2011). Thus, Sketch-based image retrieval, being a far more expressive way of image search, either alone or in conjunction with other retrieval mechanisms such as text, may yield better results, which makes it an important and interesting problem to study.

Apart from online web-scale image retrieval, an efficient sketch-based image retrieval mechanism has numerous other applications. For instance, it can be used to efficiently retrieve intended images from a constrained dataset, viz. a user's personal photo-album in a touch-sensitive camera/tablet for which no text/tag information is available. Observing the expressive power of free-hand user sketches, a few methods have been proposed for searching/designing apparels (King and Lau, 1996; Tseng *et al.*, 2009; Zhang *et al.*, 2012; Kondo *et al.*, 2014), accessories (Zeng *et al.*, 2014) or generic 3D objects such as home appliances (Eitz *et al.*, 2012b) using user-sketches. These applications require efficient shape representation and fast sketch-to-image matching to facilitate a smooth user experience. Furthermore, sketch-based retrieval can also be used to improve existing text-based image retrieval systems. For instance, it may be possible to build a sketch in an on-line manner using the first few results of a text query system (Lee and Grauman, 2009; Bagon *et al.*, 2010; Marvaniya *et al.*, 2012) and use this sketch for retrieving images that may not have any associated tag information. Image tag information may also be improved for a database in an off-line process using sketch-based retrieval.

Sketch-based Object Detection: Several approaches have been considered in the literature for describing shapes and measuring their similarity. Basic approaches for measuring the similarity of rigid shapes include Chamfer Matching (Barrow *et al.*, 1977), in which the binary template of a target shape is efficiently matched in an edge-image by calculating the distances to the nearest edge pixels. This process can be sped up using a pre-computed Distance Transform of such an edge-image. Huttenlocher *et al.* (1993) used a closely related point correspondence measure called the

[1] https://images.google.com/
[2] https://www.tineye.com/

Hausdorff distance which examines the fraction of points in one set that lie within some ε distance of points in the other set (and vice versa). Liu *et al.* (2010) improved the accuracy and efficiency of edgemap alignment substantially by incorporating image edge-orientation information and using a three-dimensional distance transform to match over possible locations and orientations.

To handle non-rigidity of shapes and speed up matching, Belongie *et al.* (2002) proposed *Shape Context* in which a shape is described by a set of descriptors, each of which captures the spatial distributions of other points along the shape contour in a log-polar space. Ling and Jacobs (2007) extend *Shape Context* by using the *inner-distance* instead of the Euclidean distance which restricts the paths between any two contour points to remain within the shape, thus making the descriptors quite robust to articulations. Gopalan *et al.* (2010) further normalize each part affinely before computing the innner distance for achieving even greater invariance to shape variations, especially in the case of perspective distortions. Shape Context-based methods have shown a good promise for clutter free shape-to-shape matchings. However, for matching shapes in images where an image has lots of extra and/or clutter edges, these are difficult to apply.

Many methods have been proposed in the literature for more sophisticated sketch-based Object Detection/Retrieval in images that handle more shape variations, although the running time for detection and retrieval is often compromised. Felzenszwalb and Schwartz (2007) use a *shape-tree* to form a hierarchical structure of contour segments for representing each object, which helps in capturing local geometric properties along with the global shape structure. Deformation is allowed at individual nodes of the tree and an efficient Dynamic Programming-based matching algorithm is used to match two curves. Ferrari *et al.* (2006) create a Contour Segment Network (CSN) by connecting nearby edge pixels in an edge-map. For matching a sketch to an image, they find paths through the constructed CSN that resemble the outline of the sketch. In a later approach, Ferrari *et al.* (2008) propose groups of k (typically ≤ 4) Adjacent approximately straight Segments (kAS) as a local feature and describe them in a scale insensitive way. These kASs are extracted for a large number of image windows and finally the object boundary is traced by linking individual small matched kASs in a multi-scale detection phase. In further work, they learn a codebook of Pairs of Adjacent Segments (PAS) (Ferrari *et al.*, 2010) which is used in combination with Hough-based centroid voting and non-rigid thin-plate spline matching for detecting the sketched object in cluttered images. In contrast to kAS matching, Ravishankar *et al.* (2008) propose a multi-stage contour-based detection approach, where Dynamic Programming is used to match segments directly to the edge pixels.

Recently, several approaches have shown a good Object Detection performance using self-contained angle-based descriptors for representing the object contours. Lu *et al.* (2009) propose a shape descriptor based on a three-dimensional histogram of angles and distances for triples of consecutive sample points along the object contours. To explicitly handle varying local shape distortions, they exploit a particle filter framework to jointly solve the contour fragment grouping and matching problems. Riemenschneider *et al.* (2010) sample a contour into a fixed number of points and calculate the angles between a line connecting any two sampled points and a line to a third point relative to the position of the first two points. This representation is insen-

sitive to translation and rotation but not scale. They employ a partial matching mechanism between two such angle-based descriptors by efficiently choosing the range of consecutive points using an integral image-based approach. To achieve robustness with respect to scale, object detection is performed over a range of scales.

All of these methods and many other state-of-the-art methods for Object Detection and Retrieval (Scott and Nowak, 2006; Kokkinos and Yuille, 2011; Ma and Latecki, 2011; Yarlagadda and Ommer, 2012) employ expensive online matching operations based on complex shape features to enhance the detection performance and typically show results on relatively small-sized datasets such as ETHZ (Ferrari *et al.*, 2006) and MPEG-7 (Latecki *et al.*, 2000) containing only a few hundreds or thousands of images while taking a considerable time to parse through each image at detection/search time. However, for a dataset with millions of images with a desired retrieval time of at most a few seconds, these methods are inapplicable/insufficient. Efficient Image pre-processing and a mechanism for fast online retrieval are necessary for large (web)-scale Image Retrieval.

Large-Scale Image Retrieval using Sketches: Only a few attempts exist in the literature for the problem of sketch-based image retrieval on large databases. Eitz *et al.* (2009) decompose an image or sketch into different spatial regions and measure the correlation between the direction of strokes in the sketch and the direction of gradients in the image by proposing two types of descriptors, viz. an Edge Histogram Descriptor and a Tensor Descriptor. Histograms of prominent gradient orientations are encoded in the Edge Histogram Descriptor, whereas the Tensor Descriptor determines a single representative vector per cell that captures the main orientation of the image gradients of that cell. Descriptors at corresponding positions in the sketch and the image are correlated for matching. Due to this strong spatial assumption, they fail to retrieve images if the sketched object is present at a different scale, orientation and/or position. To determine similar images corresponding to a sketch, a linear scan over all database images is performed. This further limits the scalability of the method for large databases.

To address the issue of scalability, Cao *et al.* (2011) propose an indexing-friendly raw contour based method. Given a sketch query, the primary objective of this method is to retrieve database images which closely correlate with the shape and the position of the sketched object. For every possible image location and a few orientations, they generate an inverted list of images that have edge pixels (edgels) at that particular location and orientation. For a sketch query, similar images are determined by counting the number of similar edgels in both the sketch and images, which makes this method susceptible to scale, translation and rotation changes. Bozas and Izquierdo (2012) introduce a hashing based framework with a strong assumption that a user only wants spatially consistent images as the search result. They extract HoG features (Dalal and Triggs, 2005) for overlapping spatial patches in an image and represent them using binary vectors by thresholding the HoG responses. The similarity between corresponding patches in the sketch and the image is estimated using the *Min-hash* algorithm (Chum *et al.*, 2008) that exploits the set-overlap similarity of these binarized descriptors. Similar to Cao *et al.* (2011), a reverse indexing structure on the hash keys is built to facilitate fast retrieval.

In order to encapsulate the spatial structure, Hu *et al.* (2010) describe both images and sketches using *Gradient field HoG* (GF-HOG) which encodes a sparse orientation field computed from the gradients of the edge pixels. To facilitate retrieval, a Bag of Visual Words model is used and sketch-to-image similarity is measured by computing the distance between corresponding frequency histograms representing the distribution of GF-HOG derived "visual words". However, this representation is noisy in presence of even small amount of background clutter. Moreover, since the experiments were performed only on small datasets that contain less than 400 images, the usability of the method on large scale retrieval is not very clear.

Riemenschneider *et al.* (2011) extend their prior idea (Riemenschneider *et al.*, 2010) to large scale retrieval by building a *vocabulary tree* (Nister and Stewenius, 2006) on the descriptors. They extract heavily overlapping contour fragments (allowing only one point shift) from a sketch and the edge-map of an image. In their approach, each contour is composed of a fixed number of points L and described as a matrix of $\binom{L}{2}$ angles that denote the orientation of the lines joining any two such points with respect to a vertical line. This makes their method sensitive to scale and orientation changes. Furthermore, due to the use of dense descriptors, the computational complexity is still very high for large datasets. Moreover, the experiments were performed only on very small datasets consisting of less than 300 images and therefore the applicability of this method in large scale is again not very clear. Zhou *et al.* (2012) determine the most "salient" object in the image and measure image similarity based on a descriptor built on the object. However, determining saliency is a very hard problem and the accuracy of even the state-of-the-art saliency methods in natural images is low (Li *et al.*, 2014), thus rendering the method possibly quite unreliable.

Proposed method in context: In this paper, we develop a system for large scale sketch-based image retrieval that can handle scale, translation and rotation (similarity) variations without compromising on efficiency, which we believe has not been addressed earlier in the literature. First, the essential shape information of all the database images is captured by extracting sequences/chains of contour segments in an offline process using two efficient and often complimentary methods: (a) finding long connected segments in contour segment networks (Sec. 2.1) and (b) using boundaries of segmented object proposals (Sec. 2.2). Such chains are represented using a similarity-invariant variable length descriptors (Sec. 3). These chain descriptors are matched using an efficient Dynamic Programming-based *approximate substring* matching algorithm. Note that, variability in the length of the descriptors makes the formulation unique and more challenging. Furthermore, partial matching is allowed to accommodate intra-class variations, small occlusions and the presence of non-object portions in the chains (Sec. 4). A hierarchical indexing tree structure of the chain descriptors of the entire image database is built offline to facilitate fast online search by matching the chains down the tree (Sec. 5). Finally, a geometric verification scheme is used for refining the retrievals by considering the geometric consistency among multiple chain matchings (Sec. 6). Results on several datasets indicate superior performance and advantages of our approach compared to prior work (Sec. 7).

2 From Images to Contour Chains

In this section, we describe offline preprocessing of database images with an objective of having a compact representation which can be used to efficiently match the images with a query sketch. We first note that a user typically draws an object along its boundary (Cole *et al.*, 2008) and a sketch of the object boundaries can more or less capture the distinctive object shape information (Cole *et al.*, 2009). Thus, an image representation based on contour information of the object boundaries would be quite appropriate in this scenario. Ferrari *et al.* (2008) construct an unweighted *Contour Segment Network* (CSN) which links nearby edge segments, and then extract $k(\leq 4)$ adjacent segments(kAS) from such a network. Full contour segment networks (Ferrari *et al.*, 2008) capture a good amount of information about the image edge segments, but are difficult to represent compactly and match efficiently in a large database. Shape Descriptors can be utilized; but they can be quite noisy, especially in the presence of image clutter. In this work, we represent shape information using (long) chains of contour segments which we believe contain a good amount of information for capturing the shape perspective and are efficient for storage and matching at the same time. While extraction of contour chains in sketches is trivial, doing so in database images is non-trivial due to the exponential number of chain possibilities. In this work, we devise two efficient and often complimentary methods for extracting and encoding the essential object boundary information in a database image by (a) finding long chains in contour segment networks, and (b) using boundaries of segmented object proposals.

2.1 Finding Long Chains in Contour Segment Network

It has been observed that long sequences of contour segments typically have a good amount of intersection with the important object boundaries. Therefore, in our approach, we try to extract long sequences of segments or chains. Furthermore, these segments must be connected to each other strongly, i.e. their connecting end points should be close to each other for them to be considered in the same chain. To this end, we extract a set of salient contours for each image that have been experimentally found to be good candidates for object boundaries and then propose a technique to group these contours into meaningful sequences that have a good chance of overlapping with the boundaries of the objects present in an image.

2.1.1 Obtaining Salient Contours

At first, in order to use fixed parameters, all the database images are normalized to a standard size (we consider the size of the longest side as 256 pixels). Then, the Berkeley Edge Detector (Martin *et al.*, 2004) is used to generate a probabilistic edge-map corresponding to the object boundaries in the image. This gives superior object boundaries compared to traditional edge detection approaches such as Canny (Canny, 1986) by considering texture along with color and brightness in an image and specifically learning for object boundaries. However such a boundary-map typically still

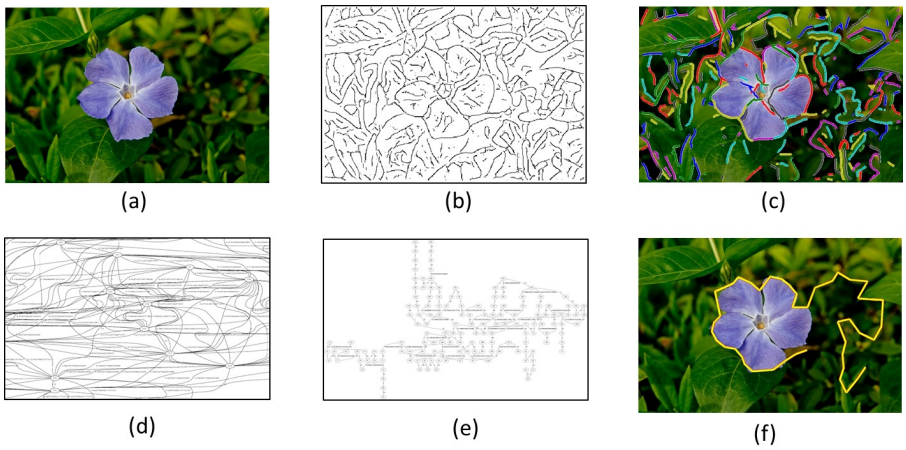

Fig. 1: Grouping Salient Contours into Chains: (a) Image, (b) Edgemap (Martin *et al.*, 2004), (c) Salient contours (Zhu *et al.*, 2007), (d) Illustrative snapshot of a constructed graph, (e) Maximum Spanning Tree for one component, (f) A long chain

contains a lot of clutter (Fig. 1(b)). Therefore, an intelligent grouping of edge pixels is done to yield better contours that have a higher chance of belonging to an object boundary. The method proposed by Zhu *et al.* (2007) groups edge pixels by considering long connected edge sequences that have as little bends as possible, especially at the junction points as it has been found that the object boundaries mostly follow a straight path at the junction points. Contours that satisfy such a constraint are called *salient* contours in their work and this method is used by us to extract a set of *salient* contours from a database image (Fig. 1(c)).

2.1.2 Segmenting Salient Contours

Given the salient contours of an image, we consider bends in them (which remain since the object shape has bends). Some articulation should be allowed at such bends since it has been observed that the object shape perspective remains relatively unchanged under articulations at such bend points (Basri *et al.*, 1998). These bend points along the contour are determined as the local maxima of the curvature. Although curvature has been defined in the literature in many ways (Cohen *et al.*, 1992; Basri *et al.*, 1998; Chetverikov, 2003), we use a simple formulation that is fast and robust. The curvature of a point p_c is obtained using m points on either side of it as:

$$\kappa_{p_c} = \sum_{i=1}^{m} w_i \cdot \angle p_{c-i} p_c p_{c+i} \tag{1}$$

where w_i is the weight defined by a Gaussian function centered at p_c. This function robustly estimates the curvature at point p_c at a given scale m. The salient contours thus obtained are split into multiple segments at such high curvature points and as a result, a set of straight line-like segments are obtained for an image.

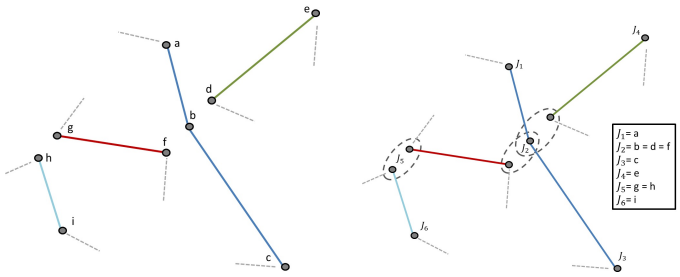

(a) (b)

Fig. 2: Graph of Joints: (a) Set of segments after breaking the salient contours (shown in different colors) at high curvature points (e.g joint b is a high curvature point). (b) Nearby endpoints of two different segments are merged if they are sufficiently close and a graph of joints (j) is created.

2.1.3 Chaining the Segments

Given a set of straight line-like segments in an image, we try to connect them. The connectivity among the segments suggests an underlying graph structure. Thus, a weighted graph/contour segment network is constructed where each end of a contour segment is considered as a vertex/joint and the edge weight between any two vertices is equal to the length of segment. Vertices from two different contour segments are merged if they are spatially close. Fig. 2(b) shows the graph corresponding to the illustrative set of straight line-like segments in Fig. 2(a).

The weight of an edge in the graph/contour segment network represents the spatial extent of the segments. Therefore, a long path in the graph based on the edge weights (higher weight is better) relates to a long connected sequence of contour segments or chains in an image. As the graph may contain cycles, to get non-cyclical paths, the maximum spanning tree[3] is constructed for each connected component in the contour segment network (Fig. 1(e)) using a standard minimum spanning tree algorithm (Cormen *et al.*, 2009) and paths are extracted from these trees. Note that the Maximum Spanning Tree algorithm removes the minimum weight edges in any cycle (Fig. 3(a)). Since we consider only long chains, minimum weight edges which correspond to the small segments would not be typically picked anyway in a long chain. For instance, in Fig. 3(a), a long chain passing through the joints J_7 and J_6 follows the path $< ..., J_7, J_4, J_5, J_6, ... >$ and therefore, removal of the smallest segment $\overline{J_7 J_6}$ in the cycle does not lead to any change in the extracted long chain. Fig. 3(b) illustrates possible chains for the contour segment network in Fig. 3(a) while Fig. 1(f) shows an extracted chain for an image (Fig. 1(a)).

A single long chain from this tree typically cannot cover the distinctive portion of any object (Fig. 4(a)). Therefore, to extract informative chains, the best N_{OC} chains are determined by allowing certain overlaps among the chains. The object boundary

[3] Maximum spanning tree of a graph can be computed by negating the edge weights and computing the minimum spanning tree.

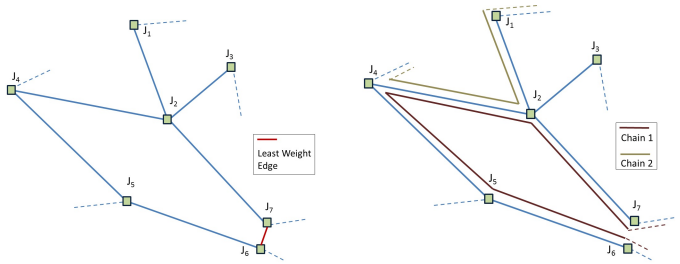

(a) (b)

Fig. 3: (a) An illustrative snapshot of a network of joints. The red colored edge joining J_6 and J_7, being the least weight edge in the induced cycle, is removed in the Maximum Spanning Tree of the network. (b) Two possible long paths/chains in the constructed tree.

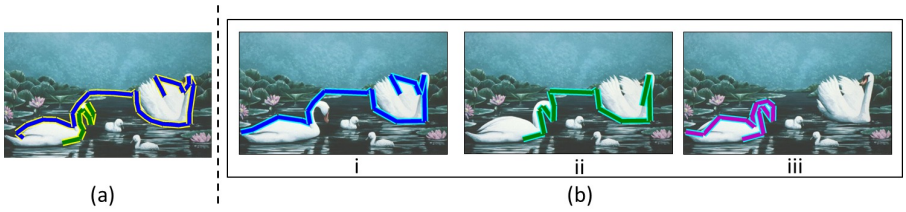

Fig. 4: (a) A single chain does not capture the object boundary, which can be addressed by considering multiple overlapping chains (b).

is typically smooth. Therefore, in the case of multiple possibilities at the junction points, we measure the smoothness of possible paths at that junction and prefer an almost straight path compared to a curve. Fig. 5 illustrates two possibilities at the junction v and the sequence $\langle ...uvw... \rangle$ is preferred to $\langle ...uvw'... \rangle$. To this end, the angle between three successive joints u, v and w is calculated. Since a straight line is preferred, the deviation of $\angle uvw$ from 180° is determined as a measure of smoothness. Furthermore, to use this measure only in the case of multiple possibilities at junction points, smoothness at a joint of a chain is normalized by the smoothness terms of all possible chains at that joint. Let t, u, v, w be four consecutive joints along a chain C and $\mathrm{d}(u, v)$ be the distance between any two joints u and v. Then the score of the chain C is determined as:

$$Score(C) = \sum_{u,v \in C} \mathrm{d}(u,v) \cdot \left[1 + \frac{\lambda_l}{2} \cdot \frac{\exp\left(-\lambda_s \cdot |\pi - \angle tuv|\right)}{\sum_{x|t,u,x \in C'} \exp\left(-\lambda_s \cdot |\pi - \angle tux|\right)} + \frac{\lambda_l}{2} \cdot \frac{\exp\left(-\lambda_s \cdot |\pi - \angle uvw|\right)}{\sum_{y|u,v,y \in C''} \exp\left(-\lambda_s \cdot |\pi - \angle uvy|\right)} \right] \tag{2}$$

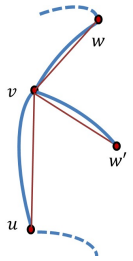

Fig. 5: $\langle ...uvw... \rangle$, being comparatively smoother, is preferred to $\langle ...uvw'... \rangle$ for scoring possible paths through the junction v

Here λ_l and λ_s are two scalar constants. C' and C'' are possible chains through the joints t, u and u, v respectively. Negative exponential functions are used since only the values close to the desired value (i.e. π) can be considered a good candidate and anything beyond a certain limit should be given a low score. Note that the smoothness term is weighted by the length of the segment in order to achieve robustness with respect to the number of intermediate joints. A tree with N_l number of leaves has $\binom{N_l}{2}$ paths and there is a substantial overlap of joints among many paths. We exploit this and use a Least Common Ancestor-based implementation (Cormen *et al.*, 2009) to efficiently score all $\binom{N_l}{2}$ paths and to sequentially select some top $N_{OC}(= 5)$ chains such that the relative overlap among them is less than $\lambda_{thresh_{chain}}(= 60\%)$. Fig. 4(b) illustrates the usefulness of considering multiple overlapping chains where only the third chain encapsulates the informative shape information of a swan. Quite reasonable results may be achieved using such long chains in a contour segment network alone (Parui and Mittal, 2014). We next propose another technique that can often provide complimentary chains in the case this algorithm fails.

2.2 Using Segmented Object Proposals

The primary objective of popular segmented object proposal techniques (Carreira and Sminchisescu, 2012; Uijlings *et al.*, 2013; Arbelaez *et al.*, 2014) is to provide locations and boundaries of the possible objects in an image. Since we need to extract boundary information for millions of images, we use a very fast method called Geodesic Object Proposals (GOP) (Krähenbühl and Koltun, 2014) for extracting a set of possible object regions. This method typically produces many overlapping scored object regions. To limit the number of chains for each image, we consider only some top $N_{GOP}(= 20)$ proposals based on their scores.

Note that while it is possible to use image segmentation techniques such as Shi and Malik (2000) and Arbelaez *et al.* (2011) to obtain good connected regions, they are typically very slow (Uijlings *et al.*, 2013). Furthermore, for the purpose of object boundary representation, the unit of interest is a single object region for each object present in an image. Therefore, Geodesic Object Proposal (Krähenbühl and Koltun, 2014) is used in this work. The method is based on superpixel growing and each proposal corresponds to a segment in the image. Thus, the object shape information

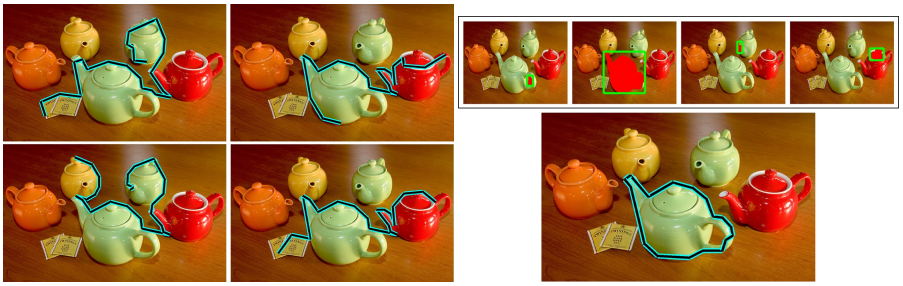

(a) (b)

Fig. 6: Limitation of considering only top N_{oc} chains from contour segment network: (a) None of the chains encloses sufficient object boundary, (b) A distinctive chain is extracted (below) considering segmented object proposals (Krähenbühl and Koltun, 2014) (above).

can be easily extracted by considering the boundaries of the proposed segments. We remove the segments that mostly touch the image boundary as such segments have incomplete boundaries and a possible object is only partially present in the image. Furthermore, to extract only distinctive object boundary, very small regions are also discarded (Fig. 6(b) and Fig. 7(a)). Finally the boundaries of the remaining regions from the top N_{GOP} proposals are taken as chains. Fig. 6(b) shows a chain successfully obtained from the segmented proposals where the chains extracted using contour segment network were inferior (Fig. 6(a)).

Fig. 7 illustrates a reverse example where the top object proposals do not contain informative object boundary information, while the overlapping chains extracted from the contour segment network covers the desired object boundary. Therefore we consider both the approaches for extracting chains for database images in an offline process.

Fig. 8 demonstrates the entire chain creation framework and Fig. 9 shows the chains thus obtained in some common images. Note that in our framework, it is easy to adapt other possibly better boundary extraction mechanisms as well due to the flexibility of the framework in terms of chain length and the number of chains.

3 Creating Descriptors for each Chain

In order to efficiently match two chains in a similarity-invariant way, we require a compact descriptor that captures the shape information of the extracted chains in a similarity-invariant way. Towards this goal, the local shape information is captured at the joints in a scale, in-plane rotation and translation invariant way. For the i^{th} joint of chain k (J_i^k), the segment length ratio $\gamma_i = \frac{l_{seg_i}}{l_{seg_{i+1}}}$ (l_{seg_i} denotes the length of the i^{th} segment) and the anti-clockwise angle θ_i (range: [0, 2π]) between the adjacent pair of segments seg_i and seg_{i+1} are determined, as shown in Fig. 10. The descriptor Ψ^k for a chain k with N segments is then defined as an ordered sequence of such

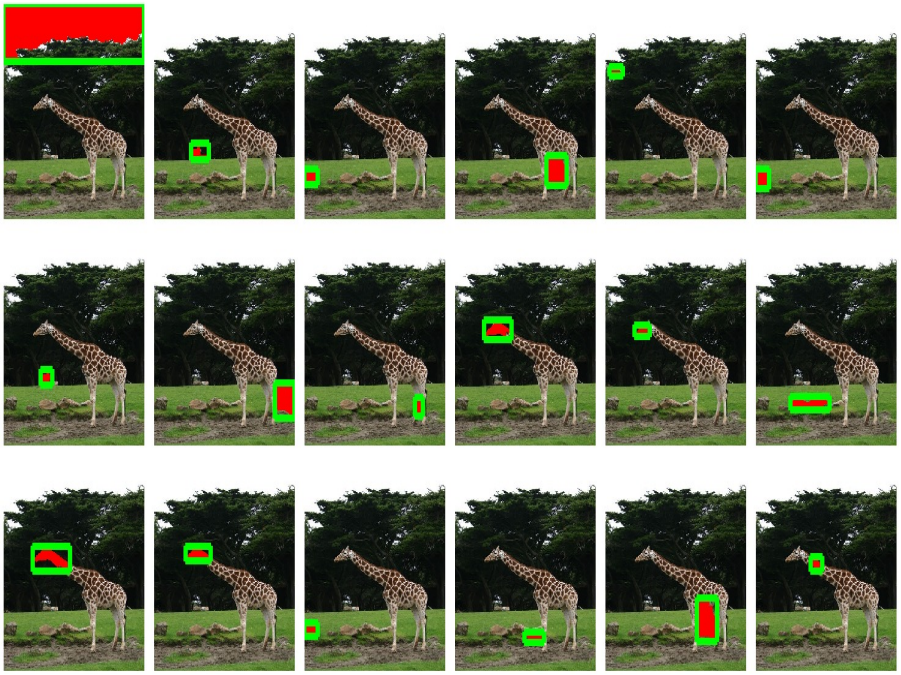

(a) Top segmented proposals obtained using the default parameters of Krähenbühl and Koltun (2014) fail to capture the object shape

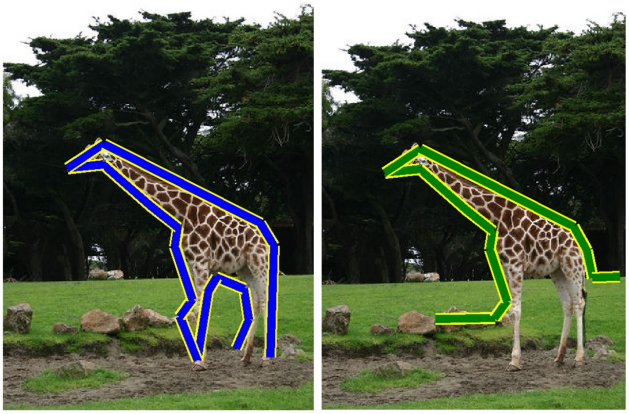

(b) Overlapping chains extracted using contour segment network cover distinctive object boundary

Fig. 7: Limitation of extracting chains only from boundaries of segmented object proposals

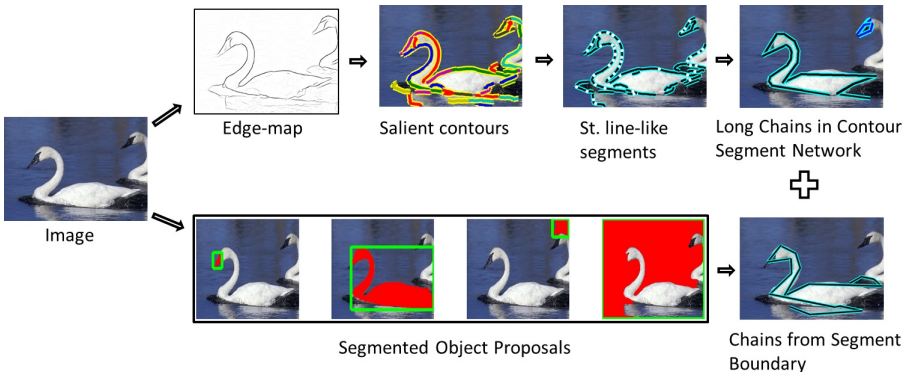

Fig. 8: Entire chain creation framework

Fig. 9: Chains extracted for some images. Different chains are represented using different colors.

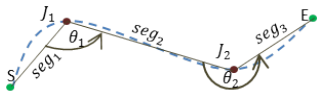

Fig. 10: The chain for the curve SE is composed of three line segments. The descriptor for this chain is $\Psi = \left\langle \gamma_i = \frac{l_{seg_i}}{l_{seg_{i+1}}}, \theta_i \mid i \in \{1,2\} \right\rangle$.

similarity-invariant quantities:

$$\Psi^k = \langle \gamma_i, \theta_i \mid i \in \{1 \ldots N-1\} \rangle \tag{3}$$

Note that Riemenschneider *et al.* (2010) also use joint information by measuring the relative angles among all pairs of sampled points along a contour. However, their representation is not scale invariant which leads to a costly online multi-scale matching phase. In contrast, our proposed descriptor is insensitive to similarities and is suitable for efficiently representing and matching contour chain information in millions of images.

Having extracted chains from images and compactly represented them in a similarity-invariant way, we next describe an approach for efficiently matching two such chains.

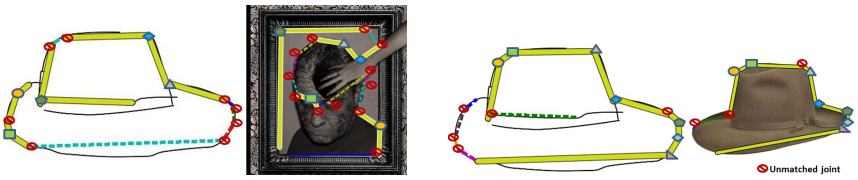

(a) (b)

Fig. 11: (a) A match when fragmented skips are allowed. (b) A match when only almost-contiguous matches are allowed. Matched joints are shown with the same marker in the sketch and the image. Unmatched portions of the chains are indicated by dashed lines.

4 Matching Two Chains

Standard vectorial type of distance measures are not applicable for matching two chains due to the variability in the lengths of the chains in our case. Further, note that the object boundary is typically captured by only a portion of the chain in the database image (Fig. 8). Therefore, a partial matching strategy of such chains needs to be devised which can be smoothly integrated with an indexing structure to efficiently determine object shape similarity.

Since image chains are typically noisy, it is not uncommon to obtain a chain that captures an object boundary and has non-object contour segments on either side of the object boundary portion. Furthermore, we assume that the object boundary is typically captured by a more or less contiguous portion of the chain without large gaps in between. Although such large split-ups may occur in certain circumstances, allowing such matches leads to a lot of false matches of images due to too much relaxation of the matching criteria. This is illustrated in Fig. 11(a), where the split matches are individually good matches but put together do not match with the intended shape structure at all. Thus, in our work, the similarity between two chains is measured by determining *the maximum (almost) contiguous matching portions of the sequences while leaving out the non-matching portions on either side from consideration* (Fig. 11(b)). This is quite similar to the Longest Common Substring[4] problem (Cormen *et al.*, 2009), with some modifications that can be solved efficiently using Dynamic Programming. We first consider the individual scores for matching two joints across two chains.

4.1 Joint Similarity

Since exact correspondence of the joints does not capture the deformation that an object may undergo, we provide a slack while matching and score the match between a pair of joints based on the deviation from exact correspondence. The score $S_{jnt}(x, y)$ for matching the x^{th} joint of chain C_1 to the y^{th} joint of chain C_2 is taken to be the

[4] Substring, unlike subsequence, does not allow gap between successive tokens.

product of two constituent scores:

$$S_{jnt}(x,y) = S_{lr}(x,y) \cdot S_{ang}(x,y) \tag{4}$$

$S_{lr}(x,y)$ is the closeness in the segment length ratio of the two adjacent segments at the x^{th} and y^{th} joints of the two descriptors:

$$S_{lr}(x,y) = \exp\left(\lambda_{lr} \cdot \left(1 - \Omega\left(\gamma_x^{C_1}, \gamma_y^{C_2}\right)\right)\right) \tag{5}$$

where $\gamma_x = \frac{l_{seg_x}}{l_{seg_{x+1}}}$ is as defined in Sec. 3, $\Omega\left(a,b\right) = \min\left(a/b, b/a\right), a, b \in \mathbb{R}_{>0}$ measures the relative similarity between two ratios ($\Omega\left(a,b\right) \in (0,1]$) and $\lambda_{lr} (= 0.5)$ is a constant. $S_{ang}(x,y)$ determines the closeness of the angles at the x^{th} and y^{th} joints and is defined as:

$$S_{ang}(x,y) = \exp\left(-\lambda_{ang} \cdot \left|\theta_x^{C_1} - \theta_y^{C_2}\right|\right) \tag{6}$$

where $\lambda_{ang}\ (= 2)$ is a constant. These two components measure the structure similarity between a pair of joints. Due to the consideration of length ratios and relative angles, the joint matching score $S_{jnt}(.,.)$ is also invariant to scale, translation and rotation.

4.2 Handling Skips

Given the scoring mechanism between a pair of joints, the match score between two chains can be determined by calculating the cumulative joint matching score of contiguous portions in the two chains. Although exact matching of such portions can be considered, due to intra-class shape variations, small partial occlusion or noise, a few non-object joints may occur in the object boundary portion of the chain. To handle these non-object portions, some skips need to be allowed. Thus, the problem is formulated as one that finds the longest *almost-contiguous* matching portion of the two chains that are to be matched. Since only descriptors are available at this stage, this matching is performed in the space of chain descriptors.

To this end, a skip penalty α is considered for the skipped joints. Note that, the loss of shape information due to a skip depends on the complexity of the skipped joints. It has been observed that a sharper angle captures more shape information than a smoother one. Hence, a skipped joint with a sharper angle should be penalized more. The sharpness (S_x) of any joint x can be calculated by taking the deviation of the joint angle (θ_x) from 180° (Fig. 12):

$$S_x = 1 - \exp(-\left|\pi - \theta_x\right|) \tag{7}$$

Furthermore, lengthier skips typically cause more loss in shape information. Therefore, to penalize skips based on its complexity, along with the sharpness of the skipped angle, the skip penalty (ω_x) is also weighted by the average length of the segments on either side of a skipped joint x:

$$\omega_x = S_x + \lambda_{skc} \cdot \frac{\left(l_{seg_x} + l_{seg_{x+1}}\right)}{2} \tag{8}$$

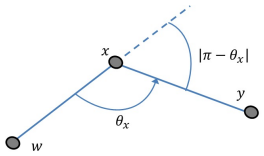

Fig. 12: Sharpness (S_x) of joint x is calculated by determining the difference of joint angle (θ_x) from π.

where, to determine the penalty of the skipped segments relative to the chain, the length of each segment is normalized by the length of the chain. λ_{skc} is a scalar constant that determines the relative effect of the two components.

4.3 Matching using Dynamic Programming

Towards finding almost-contiguous matches, one can formulate the match score $M(p_1, q_1, p_2, q_2)$ for the portion of the chain between joints p_1 and q_1 in chain C_1 and joints p_2 and q_2 in chain C_2. Let the set J_1 and J_2 denote the set of joints of chains C_1 and C_2 respectively in this interval. Also let JM be a matching between J_1 and J_2 in this interval. We restrict JM to obey the order constraint on the matches, i.e., if the joints a_1 and b_1 of the first chain are matched to the joints a_2 and b_2 respectively in the second chain, then a_1 occurring before b_1 implies that a_2 also occurs before b_2 and vice versa. Also let $X(\mathrm{JM}) = \{x|(x, y) \in \mathrm{JM}\}$ and $Y(\mathrm{JM}) = \{y|(x, y) \in \mathrm{JM}\}$ be the set of joints covered by JM. Then $M(p_1, q_1, p_2, q_2)$ is defined as:

$$M(p_1, q_1, p_2, q_2) = \max_{\substack{\mathrm{JM} \in \text{ ordered matchings in the interval } (p_1,q_1) \text{and } (p_2,q_2)}} \left(\sum_{\substack{(x,y)\in \\ \mathrm{JM}}} S_{jnt}(x, y) - \sum_{\substack{x\in \\ \mathrm{J}_1 \setminus X(\mathrm{JM})}} \omega_x^1 \alpha^1 - \sum_{\substack{y\in \\ \mathrm{J}_2 \setminus Y(\mathrm{JM})}} \omega_y^2 \alpha^2 \right) \tag{9}$$

Note that α^1 and α^2 may be different since while matching a sketch chain to an image chain, more penalty is given to a skip in the sketch chain ($\alpha = 0.07$) since it is considered cleaner and relatively more free from clutter compared to an image chain ($\alpha = 0.03$). Now, the maximum matching score ending at the joint q_1 of C_1 and q_2 of C_2 from any pair of starting joints, is defined as:

$$M(q_1, q_2) = \max_{p_1,p_2} M(p_1, q_1, p_2, q_2) \tag{10}$$

We also take the matching score of a null set ($p_1 > q_1$ or $p_2 > q_2$) as zero which constrains $M(q_1, q_2)$ to take only non-negative values. Then, it is not difficult to prove

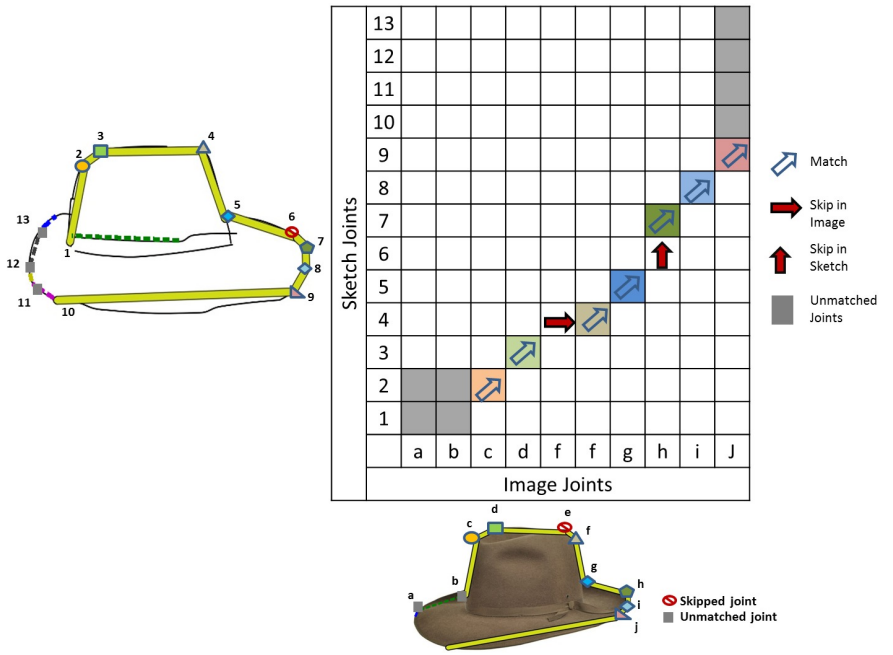

Fig. 13: Partial matching of chains with small skips. Matched joints are indicated by the same marker and colored the same in the table (Best viewed in color).

that M can be rewritten using the following recurrence relation:

$$M(q_1, q_2) = \begin{cases} 0, & \text{if } q_1, q_2 = 0 \\ \max \begin{cases} M(q_1 - 1, q_2 - 1) + S_{jnt}(q_1, q_2) \\ M(q_1 - 1, q_2) - \omega^1_{q_1} \alpha^1 \\ M(q_1, q_2 - 1) - \omega^2_{q_2} \alpha^2 \\ 0 \end{cases}, & \text{otherwise} \end{cases} \tag{11}$$

This formulation immediately leads to an efficient Dynamic Programming solution that computes M for all possible values of q_1 and q_2 starting from the first joints to the last ones. A search for the largest value of $M(q_1, q_2)$ over all possible q_1 and q_2 will then give us the best almost-contiguous matched portions between two chains C_1 and C_2 in terms of the highest matching score. Fig. 13 visually illustrates the Dynamic Programming-based matching procedure, where the chains are partially matched and a few joints are skipped while matching. This approach helps us to efficiently obtain a matching score between a pair of chains. Furthermore, to handle an object flip, we match by flipping one of the chains as well and determine the best matching score as the one that gives the highest score between the two directions. We call the final score between two chains C_1 and C_2 as the Chain Matching Score $CMS(C_1, C_2)$.

The entire operation of matching two chains takes $\mathcal{O}(n_{C_1} * n_{C_2})$ time, where n_{C_1} and n_{C_2} are the number of joints in chains C_1 and C_2 respectively. It has been

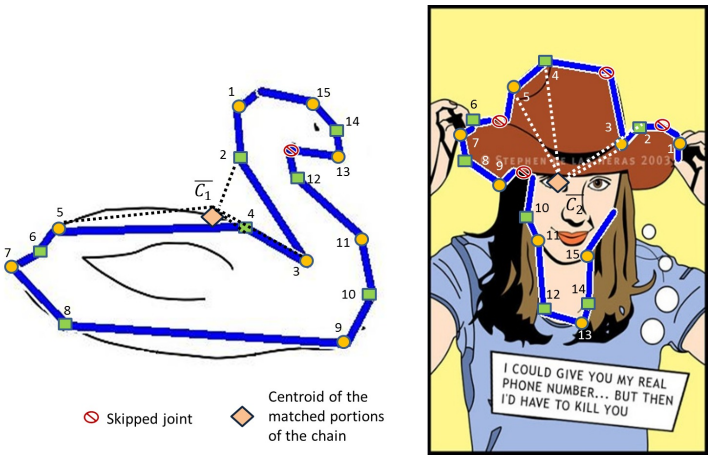

Fig. 14: Skipping of important joints can lead to a false positive match. However $\angle 2\overline{C_1}3, \angle 3\overline{C_1}4, \angle 4\overline{C_1}5$ in sketch chain and corresponding angles $\angle 2\overline{C_2}3, \angle 3\overline{C_2}4, \angle 4\overline{C_2}5$ in image chain is highly dissimilar leading to a low Global Angle Consistency score for these two falsely matched chains.

observed that a chain typically consists of 12-17 joints leading to a running time of approximately 100-400 units of joint matching, which is not very high. Note that, this DP formulation is similar to the Smith-Waterman algorithm (SW) (Smith and Waterman, 1981), which aligns two protein sequences based on a fixed alphabet-set and predefined matching costs. Meltzer and Soatto (2008) use SW to perform matching between two images under wide-baseline viewpoint changes. Our method is a slight variation from this since it performs matching based on a continuous-space formulation that measures the deviation from exact correspondence to handle deformation.

However, matching two chains by determining local joint correspondence alone sometimes leads to a globally inconsistent match as both deformation and skips of individual joints are allowed while matching. In Fig. 14, most of the joints are locally matched correctly in a similarity invariant way; but a few joints are skipped in the sketch chain and in the database chain, which leads to a globally inconsistent matching. This necessitates consideration of global consistency of the matched portions of chains for improved matching.

4.4 Global Angle Consistency of the Matched Chains

The angle that any two consecutive matched joints make with respect to some global reference point will be similar for two correctly matched chains and different for falsely-matched chains. The centroid of the matched portion of a chain is a robust point that can be used as a reference. Thus, to determine the Global Angle Consistency (GAC) between any two matched chains C_1 and C_2, we consider the centroid

of the matched portion of the chain ($\overline{C}$) as the reference point and calculate the differences in angles that any two consecutive joints make:

$$GAC(C_1, C_2) = \exp\left(-\lambda_{ac} \cdot \frac{1}{N^J} \sum_{i=1}^{N^J-1} \angle J_1^i \overline{C_1} J_1^{i+1} - \angle J_2^i \overline{C_2} J_2^{i+1}\right) \quad (12)$$

where, N^J is the total number of matched joints between C_1 and C_2 and λ_{ac} is a scalar constant. A higher value of λ_{ac} indicates a harder constraint on the global shape similarity. Fig. 14 shows an example where a false match is rejected due to a low global angle consistency score. Note that the computation of the global angle consistency is quite fast as it involves only the matched joints and can be done from the descriptor directly without referring back to the corresponding images.

The chain matching score (CMS) is weighted by the global angle consistency score (GAC) to obtain the final chain similarity score of a pair of chains $C_1$and C_2:

$$CS(C_1, C_2) = GAC(C_1, C_2) \cdot CMS(C_1, C_2) \quad (13)$$

This chain-to-chain matching strategy is used to match two image chains during indexing as well as a sketch chain to an image chain during image retrieval. Offline Image Indexing for faster retrieval is considered next.

5 Image Indexing

Given a chain descriptor, matching it online with all chains obtained from millions of images will take a considerable amount of time. Therefore, for fast retrieval of images from a large dataset, an indexing mechanism is required. Different indexing techniques have been considered in the literature for content-based image retrieval, viz. tree-based approaches using kd tree and its variants (Friedman *et al.*, 1977; Vlachos *et al.*, 2005; Muja and Lowe, 2009), hierarchical k-means (Nister and Stewenius, 2006), hashing (Indyk and Motwani, 1998; Deng *et al.*, 2011; Jegou *et al.*, 2011) etc. These approaches exploit the vectorial representation of the extracted features and perform either exact or approximate nearest neighbor search. However, in our representation, the length of each chain is not fixed. Furthermore the matching score cannot be obtained as a direct accumulation of the scores of individual dimensions. It is also not possible to use metric-based indexing techniques in our case due to a violation of the triangle inequality (Keogh and Ratanamahatana, 2005). These considerations rule out most of the possibilities such as kd tree, hashing etc. Therefore, in this work, an approach similar to hierarchical k-means (Muja and Lowe, 2009; Nister and Stewenius, 2006) but using medoids instead of means is used, which has been found to perform comparable to the state-of-the-art indexing techniques (Muja and Lowe, 2009).

All the database chains are considered for indexing and a hierarchical structure is constructed by splitting the set of extracted chains into k different clusters using the *k-medoids* algorithm (Toyama and Blake, 2002; Opelt *et al.*, 2008). At first, k chains are chosen as the cluster centroids probabilistically using the initialization mechanism of *k-means++* (Arthur and Vassilvitskii, 2007) which increases both speed and

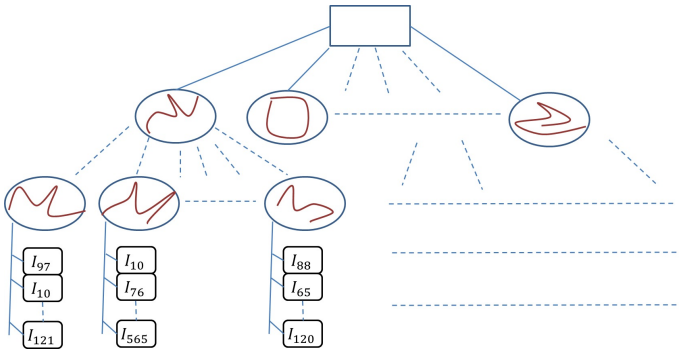

Fig. 15: Similar chains are clustered at the leaf nodes of the hierarchical k-medoid-based indexing structure.

accuracy. The remaining chains are matched to each medoid chain using Dynamic Programming-based partial matching algorithm and assigned to the closest one based on the matching score (Eqn. 13). However, due to partial matching, it is possible to get a high matching score for more than one medoids. Therefore, a chain is assigned to all the medoids for which the matching score is greater than some $Th_{ms} = (80\%)$ of the score of the closest medoid. This operation is then recursively performed on the individual clusters to determine the clusters at different levels of the tree. A leaf node of such a tree maintains a list of images of which at least one chain matches to the corresponding medoid chain. Note that, since an image has multiple chains and even one chain can belong to multiple nodes in our approach, the same image can be present at multiple leaves (Fig. 15). Given such a hierarchical chain tree constructed offline during indexing, we next discuss how to search in the Image Database given a query sketch.

6 Image Retrieval given a Query Sketch

A user typically draws an object along its boundary (Cole *et al.*, 2008). From a touch-based device, the input order of the contour points of the object boundary is usually available. Therefore, sketch chains can be trivially obtained in an online retrieval system breaking them at turns in the drawing. Offline line drawings can be decomposed in a manner similar to the edge-detected Images (Sec. 2.1.3). However, chains with less than some $Th_{nj}(= 5)$ joints are discarded as they are very simple and can match to any non-informative portion of another chain. Finally, descriptors are determined in a manner similar to image chains (Eqn. 3).

6.1 Search in the Hierarchical Chain Tree

For each of these sketch chain descriptors, a search is performed in the hierarchical k-medoids tree. At every level of the tree, the query chain is matched with all the medoid

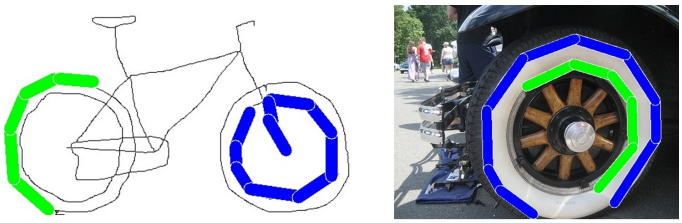

Fig. 16: Considering only individual chain matches without global consistency check can lead to a false positive retrieval.

chains and then the subtree of the best matched medoid is explored in a best-bin-first manner (Muja and Lowe, 2009). At first, a single traversal is performed through the tree following the best matched medoids at every level. This yields a small set of images corresponding to the best matched leaf medoid. Since at every level the query chain can get a good match with more than one medoids, to consider those possible matches, all the unexplored branches along the path are added to a priority queue. After the first traversal, the branch closest to the query chain is extracted from this priority queue and explored further. The search procedure stops once a pre-determined number of database images are retrieved. For all these retrieved images, at least one chain of each image matches with the query chain. Note that, for multiple sketch chains, we get multiple sets of images from the leaf nodes of the search tree, all of which are taken for the next step.

Given a set of retrieved images with corresponding matched chains, we devise a sketch-to-image matching strategy to rank the images. Since all chain matchings between a sketch and an image may not be retrieved from the hierarchical tree due to low similarity scores, we try to match the remaining chains of the sketch also with other chains of a shortlisted image to obtain the complete chain-matching information between the corresponding sketch and image. The matching score of an image for a given sketch is then calculated based on the cumulative matching scores of individual matched chain pairs between the sketch and the image. However, the actual object boundary may be split across multiple chains. Therefore it is necessary to consider all such matchings while determining the match score between a sketch and an image. However, such multiple matches may not be geometrically consistent with each other. Fig. 16 shows a case where two chains individually match well in both the sketch and the image, but the matches are not geometrically consistent with each other. This necessitates us to consider the geometric consistency of the matched chains to discard false positive retrievals. Although such geometric consistency has been studied previously in the literature (Philbin *et al.*, 2007; Sattler *et al.*, 2009; Tsai *et al.*, 2010), it is considered in a new context in this work.

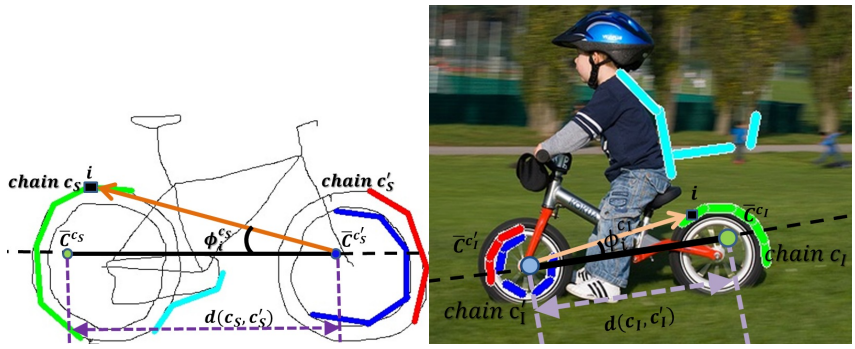

(a) (b)

Fig. 17: Pairwise geometric consistency of the matched portions of a chain pair $\mathbf{p} = (C_S, C_I)$ with respect to $\mathbf{p}' = (C'_S, C'_I)$ uses (i) the distances $d(C_S, C'_S)$ and $d(C_I, C'_I)$ between their centroids ($\overline{C}$) and (ii) the difference of angles $\left|\phi_i^{C_S} - \phi_i^{C_I}\right|$.

6.2 Geometric consistency across multiple matched chains

The geometric consistency of the matched portions of a pair of chains $\mathbf{p} = (m(C_S), m(C_I))$ with respect to that of another chain pair $\mathbf{p}' = (m(C'_S), m(C'_I))$, where C_S and C'_S are the sketch chains and C_I, C'_I are the image chains, is measured based on two factors: a) *distance-consistency* $G_d(\mathbf{p}, \mathbf{p}')$ and b) *angular-consistency* $G_a(\mathbf{p}, \mathbf{p}')$. The centroids of the matched chain portions can be obtained in a manner that is relatively robust to the presence of noise. Therefore, $G_d(\mathbf{p}, \mathbf{p}')$ is defined in terms of the closeness of the distances between the chain centroids $d(m(C_S), m(C'_S))$ in the sketch and $d(m(C_I), m(C'_I))$ in the database image (Fig. 17). These distances are normalized by the total length of the matched portions of the corresponding chains in order to achieve scale insensitivity:

$$G_d(\mathbf{p}, \mathbf{p}') = \exp\left(\lambda_c \cdot \left(1 - \Omega\left(\frac{d(m(C_S), m(C'_S))}{L_S}, \frac{d(m(C_I), m(C'_I))}{L_I}\right)\right)\right) \quad (14)$$

where, L_S=length $(m(C_S))$+ length$(m(C'_S))$, L_I=length $(m(C_I))$+ length $(m(C'_I))$, λ_c(=1) is a scalar constant and Ω is defined in Eqn. 5.

The next factor G_a measures *angular-consistency*. To achieve rotational invariance, the line joining the corresponding chain centers is considered as the reference axis and the relative angle difference at the i^{th} joint is determined (Fig. 17). $G_a(\mathbf{p}, \mathbf{p}')$ is defined using the average difference of such relative angles of all the individual matched joints in a chain:

$$G_a(\mathbf{p}, \mathbf{p}') = \exp\left(-\lambda_a \cdot \frac{1}{N^{J_\mathbf{p}}} \sum_{i=1}^{N^{J_\mathbf{p}}} \left|\phi_i^{C_S} - \phi_i^{C_I}\right|\right) \quad (15)$$

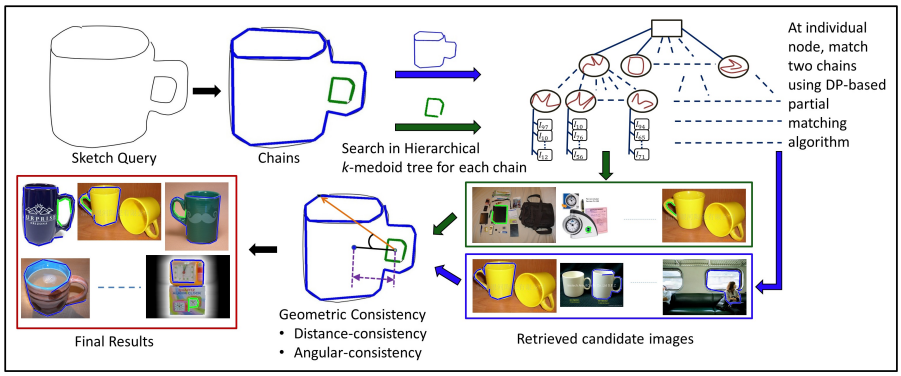

Fig. 18: The entire Retrieval Framework

where N^{J_P} is the number of matched joints between C_S and C_I and λ_a(=2) is a scalar constant. Since, both G_d and G_a should be high for consistent matching, we consider the pairwise geometric consistency $G(\mathbf{p}, \mathbf{p}')$ as a product of the constituent factors: $G(\mathbf{p}, \mathbf{p}') = G_d(\mathbf{p}, \mathbf{p}') \cdot G_a(\mathbf{p}, \mathbf{p}')$.

Erroneously matched chains are typically geometrically inconsistent with others and one may have both geometrically consistent and inconsistent pairs in a group of matched pairs between a sketch and an image. Therefore, the geometric consistency $GC(\mathbf{p})$ for a matched pair $\mathbf{p}$ is taken to be the maximum of $G(\mathbf{p}, \mathbf{p}')$ with respect to all other matched pairs $\mathbf{p}'$: $GC(\mathbf{p}) = \max_{\mathbf{p}'} G(\mathbf{p}, \mathbf{p}')$. The max operator allows us to neglect the falsely matched pairs while considering only the consistent matched pairs. Finally, the similarity score of a database image I with respect to a sketch query S is determined as:

$$Score(S, I) = \sum_{\mathbf{p} \in \mathrm{P}} GC(\mathbf{p}) \cdot CS(\mathbf{p}) \tag{16}$$

where $CS(\mathbf{p})$ is the *Chain Score* for the chain pair $\mathbf{p}$ (Eqn. 13) and P is the set of all matched pairs of chains between a sketch S and an image I. Since erroneously matched chains get very low score for consistency, effectively only the geometrically consistent chains are given weight for scoring an image. This score is used to determine the final ranking of the database images, which can be used for ranked display of such images.

Fig. 18 shows the complete retrieval framework. Results of experiments are considered next.

7 Experiments

To evaluate the performance of our system, we created a database of 1.2 million images, which contains 1 million Flickr images taken from the MIRFLICKR-1M image collection (Huiskes and Lew, 2008). In addition, we included 0.2 million images

from the Imagenet (Deng *et al.*, 2009) database in order to have some common object images in our database. In the experiments, the hierarchical index for 1.2 million images is generated with a branching factor of 32 and a maximum leaf node size of 100, which leads to a maximum tree depth of 6. Using this tree, we obtain around 1500 similar images for a given sketch, for which geometric consistency (Eqn. 16) is applied to finally rank the list of retrieved images from the chain tree.

The whole operation for a given sketch typically takes $1 - 5$ seconds on a single thread running on an Intel Core i7-3770 3.40GHz CPU. The running time typically depends on the number of chains in the sketch and most of the processing time is consumed by the geometric verification phase. However, this geometric consistency check can be trivially parallelized.

The hierarchical index for our dataset required only around 250 MB of memory. For fast online access of database chain descriptors during geometric verification and ranking of retrievals, all the descriptors for 1.2 million images are loaded a priori in the memory, which additionally required approximately 9 GB of memory. Note that, the chain descriptors can be distributed across multiple CPUs if such geometric consistency check is performed parallely. Furthermore, to make our approach work in a memory-constrained environment, for every sketch, only the descriptors corresponding to ~1500 selected images may be loaded each time in the memory at runtime although this may slow down the process somewhat due to online disk access.

Visual results for 14 sketches of different categories of varying complexity are shown in Fig. 19. These clearly indicate insensitivity of our approach to similarity transforms (e.g positive retrievals of the swan sketch). Furthermore, due to our partial matching scheme, an object is retrieved even under a viewpoint change if a portion of the distinguishing shape structure of the object is matched (e.g 4^{th} image for swan). Global invariance to similarities as well as matching with flipped objects can be seen in the results for the sketches of swan and bi-cycle (2^{nd} and 10^{th} retrieved image for swan, 2^{nd} and 9^{th} retrieved image for bi-cycle etc.). It can be easily observed that the performance of our approach depends on the complexity/distinctiveness of the shape structure. False matches (e.g cross for the sketch of airplane; two adjacent clocks/cups for the sketch of spectacle; keychain, brush for the sketch of bottle in Fig. 19) typically occur due to some shape similarity between the sketch and an object in the image, the probability of which is higher when the sketch is simple and/or contains only one chain (e.g bottle).

To understand the characteristics of the missed retrievals as well in a controlled dataset, we also tested our system on the ETHZ extended shape dataset consisting of 385 images of 7 different categories with significant scale, translation and rotation variations. Fig. 20 shows top retrieval results from the ETHZ extended shape dataset (Schindler and Suter, 2008) for a few sketches. It can be observed that the accuracy of retrieval heavily depends on the quality of the sketch. In Fig. 20, a significant portion of the first swan (7^{th} sketch) is circular and thus it matches to locally circular shapes. However, for a relatively better sketch of swan (8^{th} sketch), the number of positive retrievals is higher. Sometimes, the matched portions of two different shapes appear to be globally similar, which again leads to false positives at a few top positions (e.g matched portion of giraffe and mug for the sketch of hat in Fig. 20).

Even when the object shape is well captured in both the sketch and the image, sometimes our approach retrieves an incorrect object if the distinctive object boundary portion(s) get skipped. Fig. 21(a) illustrates such a situation where skipping of very important object boundary portions in the image of "star" leads to a wrong retrieval. Furthermore, the matched portions of corresponding chains are globally similar. Therefore, even global angular consistency check fails to identify the false retrieval in this case. This case cannot be easily addressed if we are to allow skips to handle noise in boundary extraction. Note that, for a sketch of "star", however, an image of apple does not get a high matching score because of asymmetricity in our chain matching criteria where we assume that the sketch chains are much less noisy than the image chains and hence skips are more heavily penalized in the sketch chains. This also explains the reason behind very good retrievals for the objects whose shape is (almost) unique. Fig. 21(b) also shows an example where allowing local deformation and skipping an informative joint lead to a false positive.

Quantitative measurement of the performance of a large scale retrieval system is not easy due to the difficulty in obtaining ground truth data, which is currently unavailable for a web-scale dataset. Some common metrics to measure retrieval performances (F-measure, Mean Average Precision (Manning *et al.*, 2008) etc.) use recall which is impossible to compute without a full annotation of the dataset. Therefore, to evaluate the performance of our approach quantitatively, we use the Precision-at-K measure for different rank levels (K) for the retrieval results (Rubner *et al.*, 2000). This is an acceptable measure since an end-user of a large scale Image Retrieval system typically cares only about the top results which must be good.

First, we separately evaluate the major components of our method to understand their effects. Then, we discuss the retrieval performance of the proposed algorithm in comparison to prior work on our large dataset. To study the properties of our algorithm in detail, the proposed algorithm is also evaluated on the ETHZ shape dataset (Ferrari *et al.*, 2006), on which it is possible to compute the recall as well. Finally, using these evaluations, we discuss the strengths and weaknesses of our approach.

7.1 Evaluation of major components

User sketches are highly subjective (Eitz *et al.*, 2012a) and the retrieval performance depends on the quality of the user sketch. Therefore, to obtain a robust estimate of the performance, the system must be tested using a diverse set of user sketches of varying complexity. To this end, we use a dataset (Bhattacharjee and Mittal (2015)) of 50 user-drawn sketches consisting of 10 sketches for each of the five shape categories in the ETHZ shape dataset (Ferrari *et al.*, 2006), viz. applelogo, bottle, giraffe, mug and swan. This dataset provides a wide variety in terms of the quality of the sketch and is therefore an appropriate choice for evaluation purposes. For testing these sketches, we add the images of the ETHZ shape dataset (Ferrari *et al.*, 2006) to our dataset of 1.2 million images. Finally, for each of these sketches, we determine the number of correct matches in the top 50 retrievals, where such counting has to be done manually as we do not possess any prior categorical information of the dataset images.

Table 1: Percentage of true positive images in top 50 retrievals with the corresponding standard deviation using different chain extraction mechanisms. NOC: Non-Overlapping Chains (Parui and Mittal, 2014), OC: Chains with Overlap allowed, GOP: Chains extracted using Geodesic Object Proposal (Krähenbühl and Koltun, 2014). OC+GOP uses chains extracted from both methods. GC indicates performance with geometric verification. Last row details the performance when ETHZ Models (Ferrari *et al.*, 2006) are used.

Method	**Applelogo**	**Bottle**	**Giraffe**	**Mug**	**Swan**
NOC+GC	21.4 ± 2.6	10.2 ± 2.4	7.8 ± 2.7	84.4 ± 3.2	50.9 ± 2.8
OC+GC	22.2 ± 2.4	13.8 ± 3.6	12.2 ± 4.4	90.6 ± 1.8	58.4 ± 3.4
GOP+GC	26.8 ± 3.5	13 ± 2.6	7.2 ± 2.3	91 ± 1.9	64 ± 2.9
OC+GOP	25.6 ± 1.8	14.6 ± 3.2	6.2 ± 1.4	70.4 ± 2.9	64.7 ± 2.1
OC+GOP+GC	27.6 ± 2.83	14.6 ± 3.2	12.4 ± 4.3	93.8 ± 1.5	65.1 ± 2.2
OC+GOP+GC (ETHZ Models)	54	16	20	94	80

Two chain extraction strategies are used in our work to capture the object shape information: a) overlapping long chains from contour segment networks and b) the boundaries of the segmented object proposals. To understand the relative benefits of these, we first measure the retrieval performance considering only a single method of chain extraction for all the database images at a time. Table 1 shows the retrieval scores for different object categories. Allowing overlaps while extracting long chains from contour segment networks (OC+GC) compared to using only non-overlapping chains (NOC+GC) (Parui and Mittal, 2014) increases the chance of covering the entire object boundary and hence gives superior results. When the chains are extracted using segmented object proposals (GOP+GC) (Krähenbühl and Koltun, 2014), improved results can be observed for a few categories. For some objects, segmentation produces better result and if the entire object is covered in a single segmented proposal, then an accurate chain corresponding to the object boundary can be extracted. Therefore, for such categories, viz. appelogo, swan etc., one obtains a good accuracy when chains extracted using segmented proposals are used. However, for other categories, the top proposals cover only small and/or ambiguous portions of the object(s). In such cases (giraffe, bottle), a better retrieval score is obtained when chains are extracted from the contour segment networks. To utilize the benefits of both these mechanisms, chains extracted from these two methods are combined for all the database images (OC+GOP+GC) and significant improvement on retrieval performance can be observed as compared to the different ideas considered in isolation or compared to only using the non-overlapping chains (Parui and Mittal, 2014).

Table 1 also shows the importance of considering geometric consistency of the matched chains (OC+GOP+GC vs. OC+GOP). Although significant performance gain is observed after applying geometric consistency, it is not possible to apply this step when only one chain is matched between a sketch and a database image. Due to this, for some of the categories (viz. bottle, swan etc.) in Table 1, applying geometric consistency does not make much difference in the retrieval score.

Object categories vary in the complexity and uniqueness of their shape information. For highly deformable objects,viz. giraffe, the performance is poor since our

approach can only handle a similarity transform. Furthermore, there is a considerable amount of texture variation and background clutter for the giraffe images in the ETHZ shape dataset (Ferrari *et al.*, 2006), making it hard to extract good chains for this category. For simpler shapes, viz. apple, bottle etc., many false positives get a good matching score as these shapes are relatively simpler and thus easy to match. Typically, our approach performs well for object categories with a distinctive shape and where the chains can be extracted easily. The influence of the sketch quality is also evident from the standard deviation in the retrieval accuracies. Table 1 also lists the performance of our approach for different categories using the fairly good quality ETHZ dataset model sketches (Ferrari *et al.*, 2006), for which much better performance was obtained.

Next, we show much more comprehensive retrieval results of the proposed algorithm and compare them to prior work.

7.2 Comparisons with Prior Work

7.2.1 Million Image Dataset

To evaluate our system for large scale retrieval on different object categories, we asked 5 random subjects to draw sketches for a variety of objects on a touch-based tablet and collected 75 sketches. These sketches, along with 100 sketches from a crowd-sourced sketch database (Eitz *et al.*, 2012a), containing 24 different categories in total, are used for retrieval. Non-availability of a public implementation of any prior work makes it difficult to have a comparative study with prior work. Even though a Windows phone App (SketchMatch) based on (Cao *et al.*, 2011) is available, the database is not available to make a fair comparison to other algorithms. Hence, we re-implemented this algorithm (Cao *et al.*, 2011) (EI) as well as another by Eitz *et al.* (2009) (TENSOR) and tested their algorithms on our database for the purpose of comparison. Zhou *et al.* (2012) did not provide complete implementation details in their publication and it is not trivial to make the method proposed by Riemenschneider *et al.* (2011) run efficiently on a very large database. Furthermore, Riemenschneider *et al.* (2011) did not show any result on a large scale dataset and Zhou *et al.* (2012) showed results only for 3 sketches. Hence, these methods were not compared against.

Table 2 shows the performances of our algorithm in comparison with TENSOR (Eitz *et al.*, 2009) and EI (Cao *et al.*, 2011) at different retrieval levels. First, the precision is computed for all sketches of a given object category and then the best, worst and average retrieval scores of different categories are averaged over all 24 categories. The significant deviation between the best and the worst retrieval performances indicate the diversity in the quality of the user sketches and the system response to it. It can be observed from Table 2 that our method significantly outperforms the other two methods on this large dataset. Both TENSOR (Eitz *et al.*, 2009) and EI (Cao *et al.*, 2011) consider edge matchings approximately at the same location in an image as that of the sketch and therefore, the retrieved images from their system contain the sketched shape only at the same position, scale and orientation while images containing the sketched object at a different scale, orientation and/or position are missed leading

Table 2: Precision (expressed as Percentage of true positives) at different ranks for 175 retrieval tasks in 24 categories on a dataset of 1.2 million images. B: Best, W: Worst, A: Average performances are computed among sketches for each category and then averaged.

Method	**Top 5**			**Top 10**			**Top 25**			**Top 50**			**Top 100**			**Top 250**		
	B	W	A	B	W	A	B	W	A	B	W	A	B	W	A	B	W	A
TENSOR (Eitz *et al.*, 2009)	30.8	7.5	14.7	30	7.1	13.7	24.8	7	12.9	20.8	7	12.3	16.5	5.8	10.2	9.4	3	5.7
EI (Cao *et al.*, 2011)	36.7	20.8	23.4	34.2	17.9	21.5	30	15.3	19.5	27	13.8	17.5	22.2	11.2	14.8	15.7	7.8	10.5
NOC+GC (Parui and Mittal, 2014)	80.8	42.5	60.8	72.5	38.3	53.6	54.7	29.3	39.5	40.3	20.7	28.5	31.8	16.3	22.2	23	12.5	16.5
OC+GOP+GC (Ours)	**86.5**	**43.7**	**65.5**	**78.1**	**38.4**	**58.2**	**62.4**	**30.5**	**44.6**	**51.7**	**23.8**	**34.6**	**46.7**	**21.5**	**30.8**	**34.1**	**15**	**21.7**

to false retrievals in top few matches (Fig. 22). Similar performance was observed by us on the *Sketch Match* app (SketchMatch), although a direct comparison with it is inappropriate since the databases are different. It can also be observed that using two types of chain extraction strategy and considering global angular consistency improved the performance compared to only using the non-overlapping chains (Parui and Mittal, 2014). Note that, due to the non-availability of a fully annotated dataset of a million images, it is extremely hard to use an automated parameter learning algorithm. Hence, parameters are chosen empirically by trying out a few variations. Better parameter learning/tuning could possibly improve the results further.

7.2.2 ETHZ Dataset

To provide comparisons on a standard dataset and study the recall characteristics which is difficult for a large dataset, we tested our system on the ETHZ shape dataset (Ferrari *et al.*, 2006). We used the user-drawn sketches of Bhattacharjee and Mittal (2015) for evaluation purposes. Although standard sketch-to-image matching algorithms for Object Detection that perform time consuming online processing certainly perform better than our approach on this small dataset, such comparison would be unfair since the objectives are different. Hence, we compare only against TENSOR (Eitz *et al.*, 2009) and EI (Cao *et al.*, 2011). In this dataset, we measure the percentage of positive retrievals in top 20 retrieved results which also gives an idea of recall of various approaches since the number of true positives is fixed. Table 3 shows the best, worst and average performance for the different sketches in a category (as for the previous dataset). It can be seen that our method performs much better than other methods on this dataset as well. The retrieval performance on ETHZ models (Ferrari *et al.*, 2006) further indicates substantial advantage of using very good sketches.

8 Conclusions

We have proposed an efficient image retrieval approach via hand-drawn sketches for large datasets. To the best of our knowledge, this is the first major work in the field of large scale sketch-based image retrieval that handles rotation, translation, scale

Table 3: Comparison of Percentage of true positive retrievals in top 20 using 50 sketches of Bhattacharjee and Mittal (2015) and ETHZ models (Ferrari *et al.*, 2006) on ETHZ dataset (Ferrari *et al.*, 2006).

Method	**Our Sketches**			**ETHZ Models**
	Best	Worst	Avg	(Ferrari *et al.*, 2006)
TENSOR (Eitz *et al.*, 2009)	17	10	13.5	13.6
EI (Cao *et al.*, 2011)	46	6	26.7	27.9
NOC+GC (Parui and Mittal, 2014)	60	28	42.9	49.3
OC+GOP+GC (Ours)	**76**	**37**	**56.3**	**76.4**

and small variations of the object shape even for a dataset consisting of millions of images. This is accomplished by representing the images using chains of contour segments that have a high probability of containing the object boundary. A similarity-invariant variable length descriptor is proposed that is used to partially match two chains in a hierarchical indexing framework. We also proposed a geometric verification scheme for improving the search accuracy. Experimental results shown on different datasets clearly indicate the benefits of our approach compared to the existing methods.

Due to similarity-invariance of our approach as compared to other relevant work, our method could be used to efficiently search in large natural image databases, which typically have a lot of variations. The proposed method could also open the window for efficiently searching in constrained image databases, viz. personal photo albums, which typically do not contain any tag/text information. Furthermore, our method could be augmented by other techniques such as text for tagging images in an offline fashion.

One major issue of our approach is the difficulty in extracting "good" representative chains in the presence of considerable background clutter. Newer and better boundary extraction mechanisms can be easily adapted to our framework for improving the quality of the chains. Furthermore, the proposed sketch-to-image matching approach is only similarity-invariant and also fails to handle substantial deformation and major viewpoint changes. A sophisticated affine or even projective-invariant matching mechanism could possibly help retrieve images with such variations as well and can be considered in future work.

References

Arbelaez P, Maire M, Fowlkes C, Malik J (2011) Contour detection and hierarchical image segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence 33(5):898–916

Arbelaez P, Pont-Tuset J, Barron J, Marques F, Malik J (2014) Multiscale combinatorial grouping. In: IEEE International Conference on Computer Vision and Pattern Recognition

Arthur D, Vassilvitskii S (2007) k-means++: The advantages of careful seeding. In: ACM-SIAM Symposium on Discrete algorithms

Bagon S, Brostovski O, Galun M, Irani M (2010) Detecting and sketching the common. In: IEEE International Conference on Computer Vision and Pattern Recognition

Barrow H, Tenenbaum J, Bolles R, Wolf H (1977) Parametric correspondence and chamfer matching: two new techniques for image matching. In: International Joint Conference on Artificial Intelligence

Basri R, Costa L, Geiger D, Jacobs D (1998) Determining the similarity of deformable shapes. Vision Research 38(15):2365–2385

Belongie S, Malik J, Puzicha J (2002) Shape matching and object recognition using shape contexts. IEEE Transactions on Pattern Analysis and Machine Intelligence 24(4):509–522

Bhattacharjee SD, Mittal A (2015) Part-based Deformable Object Detection with a Single Sketch (SUBMITTED to Computer Vision and Image Understanding)

Bozas K, Izquierdo E (2012) Large scale sketch based image retrieval using patch hashing. In: Advances in Visual Computing

Canny J (1986) A computational approach to edge detection. IEEE Transactions on Pattern Analysis and Machine Intelligence (6):679–698

Cao Y, Wang C, Zhang L, Zhang L (2011) Edgel index for large-scale sketch-based image search. In: IEEE International Conference on Computer Vision and Pattern Recognition

Carreira J, Sminchisescu C (2012) CPMC: Automatic object segmentation using constrained parametric min-cuts. IEEE Transactions on Pattern Analysis and Machine Intelligence 34(7):1312–1328

Chetverikov D (2003) A simple and efficient algorithm for detection of high curvature points in planar curves. In: Computer Analysis of Images and Patterns

Chum O, Philbin J, Zisserman A (2008) Near Duplicate Image Detection: min-Hash and tf-idf Weighting. In: British Machine Vision Conference

Cohen I, Ayache N, Sulger P (1992) Tracking Points on Deformable Objects Using Curvature Information. In: European Conference on Computer Vision

Cole F, Golovinskiy A, Limpaecher A, Barros HS, Finkelstein A, Funkhouser T, Rusinkiewicz S (2008) Where do people draw lines? In: ACM Transactions on Graphics (Proceedings of SIGGRAPH)

Cole F, Sanik K, DeCarlo D, Finkelstein A, Funkhouser T, Rusinkiewicz S, Singh M (2009) How well do line drawings depict shape? In: ACM Transactions on Graphics (Proceedings of SIGGRAPH)

Cormen TH, Leiserson CE, Rivest RL, Stein C (2009) Introduction to Algorithms, 3rd edn. The MIT Press

Dalal N, Triggs B (2005) Histograms of oriented gradients for human detection. In: IEEE International Conference on Computer Vision and Pattern Recognition

Deng J, Dong W, Socher R, Li LJ, Li K, Fei-Fei L (2009) Imagenet: A large-scale hierarchical image database. In: IEEE International Conference on Computer Vision and Pattern Recognition

Deng J, Berg AC, Fei-Fei L (2011) Hierarchical semantic indexing for large scale image retrieval. In: IEEE International Conference on Computer Vision and Pattern Recognition

Eitz M, Hildebrand K, Boubekeur T, Alexa M (2009) A descriptor for large scale image retrieval based on sketched feature lines. In: Eurographics Symposium on Sketch-Based Interfaces and Modeling

Eitz M, Hays J, Alexa M (2012a) How do humans sketch objects? In: ACM Transactions on Graphics (Proceedings of SIGGRAPH)

Eitz M, Richter R, Boubekeur T, Hildebrand K, Alexa M (2012b) Sketch-Based Shape Retrieval. In: ACM Transactions on Graphics (Proceedings of SIGGRAPH)

Felzenszwalb PF, Schwartz JD (2007) Hierarchical matching of deformable shapes. In: IEEE International Conference on Computer Vision and Pattern Recognition

Ferrari V, Tuytelaars T, Van Gool L (2006) Object detection by contour segment networks. In: European Conference on Computer Vision

Ferrari V, Fevrier L, Jurie F, Schmid C (2008) Groups of adjacent contour segments for object detection. IEEE Transactions on Pattern Analysis and Machine Intelligence 30(1):36–51

Ferrari V, Jurie F, Schmid C (2010) From images to shape models for object detection. International Journal of Computer Vision 87(3):284–303

Friedman JH, Bentley JL, Finkel RA (1977) An algorithm for finding best matches in logarithmic expected time. ACM Transactions on Mathematical Software 3(3):209–226

Gopalan R, Turaga P, Chellappa R (2010) Articulation-invariant representation of non-planar shapes. In: European Conference on Computer Vision

Hu R, Barnard M, Collomosse J (2010) Gradient field descriptor for sketch based retrieval and localization. In: International Conference on Image Processing

Huiskes MJ, Lew MS (2008) The MIR flickr retrieval evaluation. In: Proceedings of the 1st ACM International Conference on Multimedia Information Retrieval

Huttenlocher DP, Klanderman GA, Rucklidge WJ (1993) Comparing images using the Hausdorff distance. IEEE Transactions on Pattern Analysis and Machine Intelligence 15(9):850–863

Indyk P, Motwani R (1998) Approximate nearest neighbors: towards removing the curse of dimensionality. In: Proceedings of the 30th annual ACM Symposium on Theory of Computing

Jegou H, Douze M, Schmid C (2011) Product quantization for nearest neighbor search. IEEE Transactions on Pattern Analysis and Machine Intelligence 33(1):117–128

Keogh E, Ratanamahatana CA (2005) Exact indexing of dynamic time warping. Knowledge and Information Systems 7(3):358–386

King I, Lau TK (1996) A feature-based image retrieval database for the fashion, textile, and clothing industry in Hong Kong. In: International Symposium Multi-Technology Information Processing

Kokkinos I, Yuille A (2011) Inference and learning with hierarchical shape models. International Journal of Computer Vision 93(2):201–225

Kondo Si, Toyoura M, Mao X (2014) Sketch based skirt image retrieval. In: ACM Joint Symposium on Computational Aesthetics, Non-Photorealistic Animation and Rendering, and Sketch-Based Interfaces and Modeling

Krähenbühl P, Koltun V (2014) Geodesic object proposals. In: European Conference on Computer Vision

Landay JA, Myers BA (2001) Sketching interfaces: Toward more human interface design. IEEE Computer 34(3):56–64

Latecki LJ, Lakamper R, Eckhardt T (2000) Shape descriptors for non-rigid shapes with a single closed contour. In: IEEE International Conference on Computer Vision and Pattern Recognition

Lee YJ, Grauman K (2009) Shape discovery from unlabeled image collections. In: IEEE International Conference on Computer Vision and Pattern Recognition

Lee YJ, Zitnick CL, Cohen MF (2011) Shadowdraw: real-time user guidance for freehand drawing. In: ACM Transactions on Graphics (Proceedings of SIGGRAPH)

Li Y, Hou X, Koch C, Rehg JM, Yuille AL (2014) The secrets of salient object segmentation. In: IEEE International Conference on Computer Vision and Pattern Recognition

Ling H, Jacobs DW (2007) Shape classification using the inner-distance. IEEE Transactions on Pattern Analysis and Machine Intelligence 29(2):286–299

Liu MY, Tuzel O, Veeraraghavan A, Chellappa R (2010) Fast directional chamfer matching. In: IEEE International Conference on Computer Vision and Pattern Recognition

Lu C, Latecki LJ, Adluru N, Yang X, Ling H (2009) Shape guided contour grouping with particle filters. In: IEEE International Conference on Computer Vision

Ma T, Latecki LJ (2011) From partial shape matching through local deformation to robust global shape similarity for object detection. In: IEEE International Conference on Computer Vision and Pattern Recognition

Manning CD, Raghavan P, Schütze H (2008) Introduction to information retrieval, vol 1. Cambridge University Press

Marr D (1982) Vision: A Computational Investigation into the Human Representation and Processing of Visual Information. Henry Holt and Co. Inc.

Martin DR, Fowlkes CC, Malik J (2004) Learning to detect natural image boundaries using local brightness, color, and texture cues. IEEE Transactions on Pattern Analysis and Machine Intelligence 26(5):530–549

Marvaniya S, Bhattacharjee S, Manickavasagam V, Mittal A (2012) Drawing an automatic sketch of deformable objects using only a few images. In: European Conference on Computer Vision. Workshops and Demonstrations

Meltzer J, Soatto S (2008) Edge descriptors for robust wide-baseline correspondence. In: IEEE International Conference on Computer Vision and Pattern Recognition

Muja M, Lowe DG (2009) Fast Approximate Nearest Neighbors with Automatic Algorithm Configuration. In: International Conference on Computer Vision Theory and Application

Nister D, Stewenius H (2006) Scalable recognition with a vocabulary tree. In: IEEE International Conference on Computer Vision and Pattern Recognition

Opelt A, Pinz A, Zisserman A (2008) Learning an alphabet of shape and appearance for multi-class object detection. International Journal of Computer Vision 80(1):16–44

Parui S, Mittal A (2014) Similarity-Invariant Sketch-Based Image Retrieval in Large Databases. In: European Conference on Computer Vision

Philbin J, Chum O, Isard M, Sivic J, Zisserman A (2007) Object retrieval with large vocabularies and fast spatial matching. In: IEEE International Conference on Com-

puter Vision and Pattern Recognition
Ravishankar S, Jain A, Mittal A (2008) Multi-stage contour based detection of deformable objects. In: European Conference on Computer Vision
Riemenschneider H, Donoser M, Bischof H (2010) Using partial edge contour matches for efficient object category localization. In: European Conference on Computer Vision
Riemenschneider H, Donoser M, Bischof H (2011) Image retrieval by shape-focused sketching of objects. In: Computer Vision Winter Workshop
Rubner Y, Tomasi C, Guibas LJ (2000) The earth mover's distance as a metric for image retrieval. International Journal of Computer Vision 40(2):99–121
Sattler T, Leibe B, Kobbelt L (2009) SCRAMSAC: Improving RANSAC's efficiency with a spatial consistency filter. In: IEEE International Conference on Computer Vision and Pattern Recognition
Schindler K, Suter D (2008) Object detection by global contour shape. Pattern Recognition 41(12):3736–3748
Schroff F, Criminisi A, Zisserman A (2011) Harvesting image databases from the web. IEEE Transactions on Pattern Analysis and Machine Intelligence 33(4):754–766
Scott C, Nowak R (2006) Robust contour matching via the order-preserving assignment problem. IEEE Transactions on Image Processing 15(7):1831–1838
Shi J, Malik J (2000) Normalized cuts and image segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence 22(8):888–905
Sigurbjörnsson B, Van Zwol R (2008) Flickr tag recommendation based on collective knowledge. In: International conference on World Wide Web
SketchMatch (2011) http://research.microsoft.com/en-us/projects/sketchmatch
Smith TF, Waterman MS (1981) Identification of common molecular subsequences. Journal of molecular biology 147(1):195–197
Toyama K, Blake A (2002) Probabilistic tracking with exemplars in a metric space. International Journal of Computer Vision 48(1):9–19
Tsai SS, Chen D, Takacs G, Chandrasekhar V, Vedantham R, Grzeszczuk R, Girod B (2010) Fast geometric re-ranking for image-based retrieval. In: International Conference on Image Processing
Tseng CH, Hung SS, Tsay JJ, Tsaih D (2009) An efficient garment visual search based on shape context. WSEAS transactions on Computers 8(7):1195–1204
Uijlings JR, van de Sande KE, Gevers T, Smeulders AW (2013) Selective search for object recognition. International Journal of Computer Vision 104(2):154–171
Vlachos M, Vagena Z, Yu PS, Athitsos V (2005) Rotation invariant indexing of shapes and line drawings. In: ACM International Conference on Information and Knowledge Management
Walther DB, Chai B, Caddigan E, Beck DM, Fei-Fei L (2011) Simple line drawings suffice for functional MRI decoding of natural scene categories. Proceedings of the National Academy of Sciences 108(23):9661–9666
Yarlagadda P, Ommer B (2012) From meaningful contours to discriminative object shape. In: European Conference on Computer Vision

Zeng L, Liu Yj, Wang J, Zhang Dl, Yuen MMF (2014) Sketch2Jewelry: Semantic feature modeling for sketch-based jewelry design. Computers & Graphics 38:69–77

Zhang W, Antunez E, Gokturk S, Sumengen B (2012) Apparel silhouette attributes recognition. In: Workshop on Applications of Computer Vision

Zhou R, Chen L, Zhang L (2012) Sketch-based image retrieval on a large scale database. In: ACM international conference on Multimedia

Zhu Q, Song G, Shi J (2007) Untangling cycles for contour grouping. In: IEEE International Conference on Computer Vision

Fig. 19: Top retrieved images for 14 sketches from 1.2 million images. Retrieved images indicate similarity insensitivity and deformation handling of our approach. Chains are embedded on the retrieved images to illustrate the locations of the matchings. Multiple matched chains are shown using different colors. Correct, similar and false matches are illustrated by green, yellow and red boxes respectively (Best viewed in color).

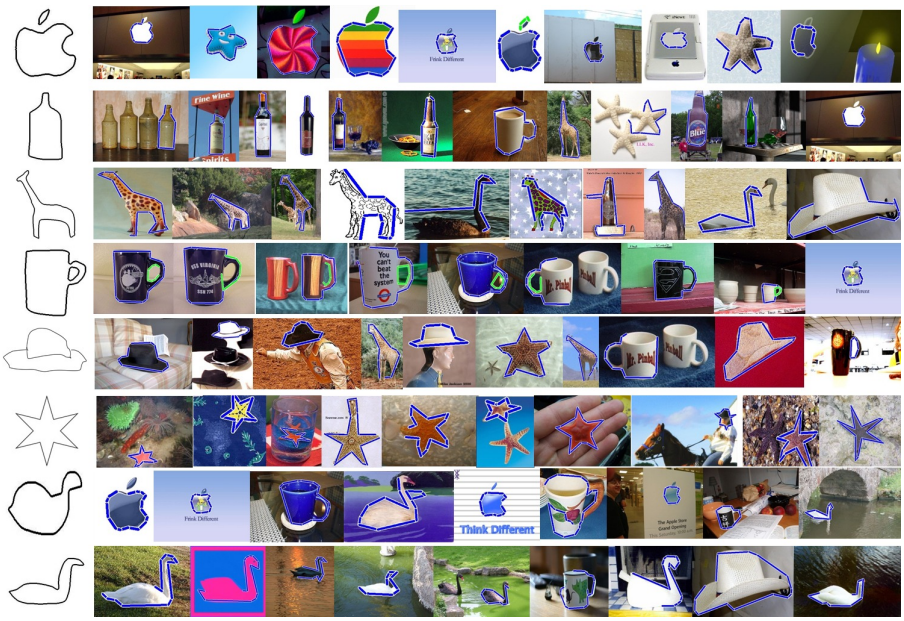

Fig. 20: Top retrievals from ETHZ extended shape dataset (Schindler and Suter, 2008) for few sketches.

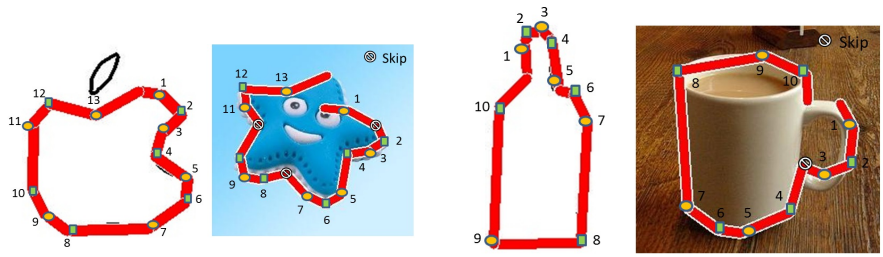

(a) (b)

Fig. 21: Skipping of important joints and allowing local deformations lead to false positive retrievals. Matched joints are numbered same in both the sketch and the image.

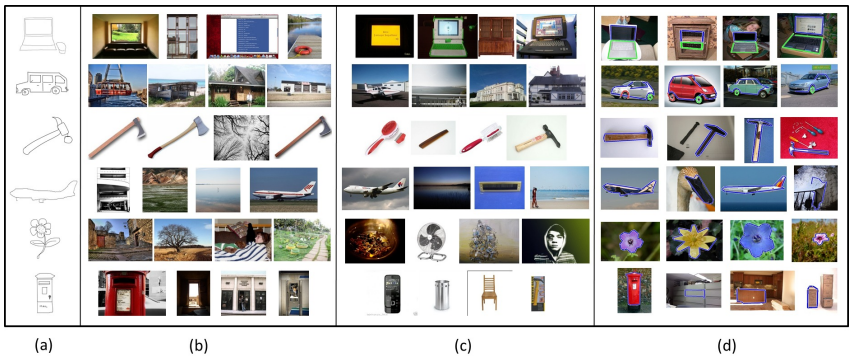

Fig. 22: Top 4 results by (b) Eitz *et al.* (2009), (c) Cao *et al.* (2011) and (d) our system on a 1.2 million image dataset for some sample sketches (a).